

Table of Contents

Front-End Usage Guidelines.....1
 Pleiades Front-End Usage Guidelines.....1
 Columbia Front-End Usage Guidelines.....3

Front-End Usage Guidelines

Pleiades Front-End Usage Guidelines

DRAFT

This article is being reviewed for completeness and technical accuracy.

The front-end systems pfe[1-12] and bridge[1,2] provide an environment that allows you to get quick turnaround while performing the following:

- file editing
- compiling
- short debugging and testing session
- batch job submission to the compute systems

Bridge[1,2], with 4 times the memory on pfe[1-12] and better interconnects, can also be used for the following two functions:

1. Post processing

These nodes have 64-bit versions of IDL, Matlab, and Tecplot installed and have 64 GB of memory (4 times the amount of memory on pfe[1-12]). The bridge nodes will run these applications much faster than on pfe[1-12].

2. File transfer between Pleiades and Columbia or Lou

Note that both the Pleiades Lustre filesystems (/nobackupp[10-70]) and the Columbia CXFS filesystems (/nobackup1[1-h], /nobackup2[a-i]) are mounted on the bridge nodes.

To copy files between the Pleiades Lustre and Columbia CXFS filesystems, log in to bridge[1,2] and use the *cp* command to perform the transfer. The 10 Gigabit Ethernet (GigE) connections on the two bridge nodes are faster than the 1 GigE used on pfe[1-12], therefore, file transfer out of Pleiades is improved when using the bridge nodes.

File transfers from bridge[1,2] to Lou[1,2] will go over the 10 GigE interface by default. The commands *scp*, *bbftp*, and *bbscp* are available to do file transfers. Since *bbscp* uses almost the same syntax as *scp*, but performs faster than *scp*, we recommend using *bbscp* over *scp* in cases where you do not require the data to be encrypted when sent over the network.

The pfe systems ([pfe1-12]) have a 1 GigE connection, which can be saturated by a single secure copy (scp). You will see bad performance whenever more than one file transfer is happening. Use of bridge1 and bridge2 for file transfers is strongly recommended.

File transfers from the compute nodes to Lou must go through pfe[1-12] or bridge[1,2] first, although going through bridge[1,2] is preferred for performance consideration. See [Transferring Files from the Pleiades Compute Nodes to Lou](#) for more information.

When sending data to Lou[1-2], please keep your largest individual file size under 1 TB, as large files will keep all of the tape drives busy, preventing other file restores and backups. To prevent the filesystems on Lou[1-2] from filling up, please limit total data transfers to 1 TB and then wait an hour before continuing. This allows the tape drives to write the data to tape.

Additional restrictions apply to using these front-end systems:

1. No MPI jobs are allowed to run on pfe[1-12], bridge[1,2]
2. A job on pfe[1-12] should not use more than 8 GB. When it does, a courtesy email is sent to the owner of the job.
3. A job on bridge[1,2] should not use more than 56 GB. When it does, a courtesy email is sent to the owner of the job.

Columbia Front-End Usage Guidelines

DRAFT

This article is being reviewed for completeness and technical accuracy.

The front-end system, cfe2, provide an environment that allows users to get quick turnaround while performing the following: file editing; file management; short debugging and testing sessions; and batch job submission to the compute systems.

Running long and/or large (in terms of memory and/or number of processors) debugging or production jobs interactively or in the background of cfe2 is considered to be inconsiderate behavior to the rest of the user community. If you need help submitting such jobs to the batch systems, please contact a NAS scientific consultant at (650) 604-4444 or 1-800-331-USER or send e-mail to: support@nas.nasa.gov

Jobs that cause significant impact on the system load of the Columbia front-end machine (cfe2) are candidates for removal in order to bring the front-end systems back to a normal and smooth environment for all users. A *cron* job regularly monitors the system load and determines if job removal is necessary. The criteria for job removal are described below. Owners of any removed jobs will receive a notification e-mail.

1. To be eligible for removal, the number of processors a front-end interactive job uses can be one (1) or more. Exceptions to this are those programs, utilities, etc. common to users and/or NASA missions that are listed in an "exception file". Examples of these would be:

bash cp csh emacs gzip rsync scp sftp sh ssh tar tcsh

Users can submit program names to be added to this exception file by mailing requests to: support@nas.nasa.gov

2. For qualifying processes, the CPU time usage of each process in a job has, on the average, exceeded a threshold defined as:

$(20 \text{ min} \times 8 / \text{number of processes for the job})$

That is, a baseline for removal is a job with 8 processors running for more than 20 minutes. The maximum amount of time allowed for each processor in a job is scaled using the formula:

$20 \text{ min} \times 8 \text{ cpu} / \text{number-of-processes}$

Therefore, the following variations are possible:

◆ $160 \text{ minutes} = (20 \times 8) / 1 \text{ cpu}$

- ◆ 80 minutes = $(20 * 8) / 2$ cpu
- ◆ 40 minutes = $(20 * 8) / 4$ cpu
- ◆ 20 minutes = $(20 * 8) / 8$ cpu
- ◆ 10 minutes = $(20 * 8) / 16$ cpu
- ◆ 5 minutes = $(20 * 8) / 32$ cpu
- ◆ 2.5 minutes = $(20 * 8) / 64$ cpu

The conditions of removal are subject to change, when necessary.